Reasons for Using Jacquard

Richard A. Gerber
NERSC User Services
November 2009

1. Overview

In the spring and summer of 2009, NERSC asked members of the top 20 projects using Jacquard "Why do you use Jacquard?" Representatives from 13 projects responded. If you would like to contribute to the survey, we welcome email to ragerber@lbl.gov.

2. Summary of User Responses to "Why do you use Jacquard?"

Reason Cited	Number
Fast turnaround	6
Jacquard has software not installed on Franklin: Gaussian, GROMACS	4
Jacquard is more economical than Franklin (or Bassi) because of 0.6 charge factor	4
Jacquard is more stable than Franklin	3
Jacquard has longer runtime limits on queues (48 hours) than Franklin	3
Jacquard has more per-core memory than Franklin	2
Increased throughput running on Jacquard and other machine simultaneously.	1
Jobs require local /tmp directories on the nodes	1
Jacquard is faster than Franklin	1
Codes call SYSTEM(); not available on Franklin	1
Effort to port code: already have code developed, debugged, and verified on Jacquard	1

3. Jacquard's Unique Features

- Full Linux OS on compute nodes.
- 3 GB memory per core (2 GB per core on Franklin)

4. Jacquard 2009 Usage Profile

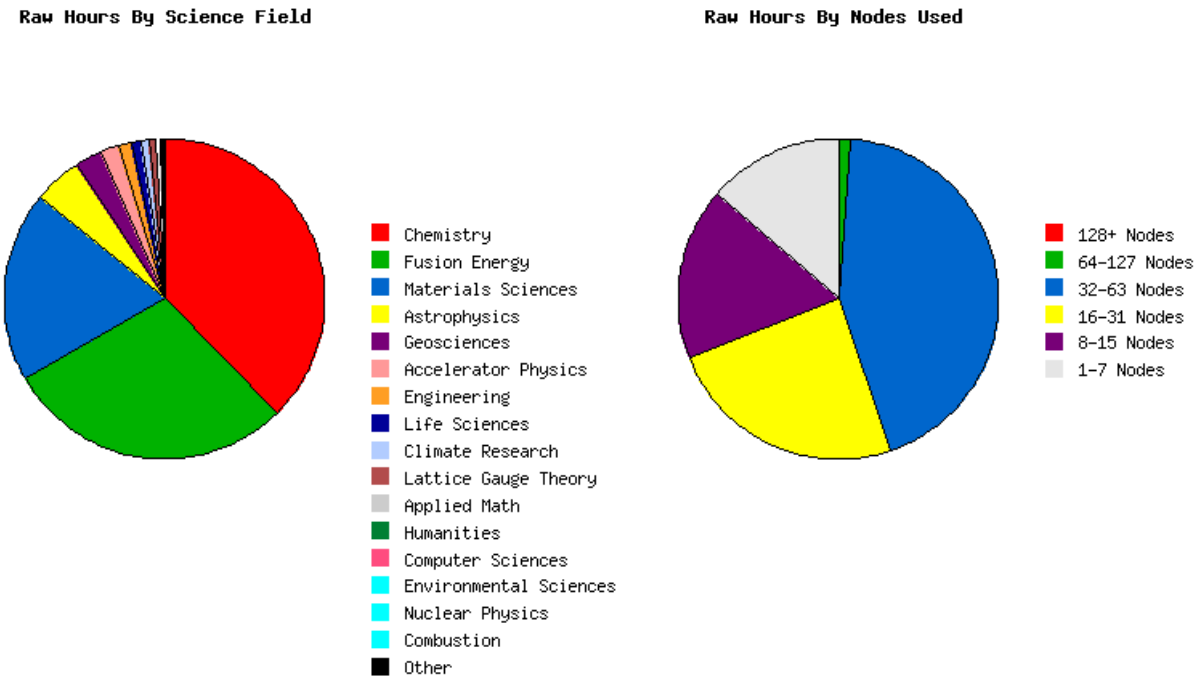

5. Top Projects Using Jacquard

Repositories with the most usage on Jacquard through the first quarter of the 2009 Allocation Year.

Repo	Project Title	PI	Science Category	Jacquard Raw Hours	Jacquard Avg. Cores	Franklin Raw Hours	Franklin Avg. Cores
m651	Theoretical studies of combustion dynamics	Joel Bowman, Emory University	Chemistry	422,309	64	201	16
m631	Hydrocarbons in plasma	Bastiaan Braams, Emory University	Fusion Energy	365,641	64	79	4
gc3	Magnetic Fusion Plasma Microturbulence Project	Bruce Cohen, Lawrence Livermore National Laboratory	Fusion Energy	325,260	32	1,146,688	1,280
m526	Computational Resources for the Nanomaterials Theory Institute at the Center for Nanophase Materials Sciences	Paul Kent, Oak Ridge National Laboratory	Materials Sciences	141,128	24	3,023,860	512
m808	SciDAC GSEP: Gyrokinetic Simulation of Energetic Particle Turbulence and Transpor	Zhihong Lin, University of California, Irvine	Fusion Energy	113,322	64	3,914,538	2,080
m93	Quantum Dynamical Study of Charge and Energy Tranfer Processes in the CondensedPhase	Haobin Wang, New Mexico State University	Chemistry	97,386	8	809,683	128
mp208	Quantum Monte Carlo for Electronic Structure of Molecules	William Lester, University of California-Berkeley	Chemistry	91,931	32	1,731,207	4,096
m881	High dimensional quantum dynamics studies of molecular spectroscopy	Hua-Gen Yu, Brookhaven National Laboratory	Chemistry	75,471	32	7,938	180
m530	Aqueous Electrolyte Solutions at Solid-Fluid Interfaces	Alberto Striolo, University of Oklahoma	Materials Sciences	66,454	16	139,067	60
m554	Calibration of 3D Upper Mantle Structure in Eurasia Using Regional and Teleseismic Full Waveform Seismic Data	Barbara Romanowicz, University of California, Berkeley	Geosciences	64,556	32	0	–

Repo	Project Title	PI	Science Category	Jacquard Raw Hours	Jacquard Avg. Cores	Franklin Raw Hours	Franklin Avg. Cores
m783	Computational Studies at BNL of the Chemistry of Energy Production and Use	James Muckerman, Brookhaven National Laboratory	Chemistry	63,350	28	39,062	32
m289	Evolution of Massive Galaxies and Black Holes	Chung-Pei Ma, UC-Berkeley	Astrophysics	60,165	32	0	–
m779	Baryon Oscillation Spectroscopic Survey	Peter Nugent, Lawrence Berkeley National Laboratory	Astrophysics	59,465	64	0	–
mp2	LLNL MFE Supercomputing	Bruce Cohen, Lawrence Livermore National Laboratory	Fusion Energy	53,924	32	0	–
m249	Computational chemistry search of efficient catalysts	Perla Balbuena, Texas A&M University	Chemistry	53,305	16	339,701	32
m622	Investigating quantum dots for high efficiency fourth generation photovoltaic	Fritz Prinz, Stanford University	Materials Sciences	51,570	32	18,825	16
m624	Multiscale Modeling of Molecular Delivery in Nanomedicine and Nanofluidics	Petr Kral, University of Illinois, Chicago	Materials Sciences	49,855	32	160,273	140
mp311	Simulations of electron-hole relaxation dynamics on functionalized TiO2-anatasesurfaces	Victor Batista, Yale University	Chemistry	43,256	64	0	–
m474	Polymers in confined geometries	Roland Faller, University of California, Davis	Materials Sciences	42,115	16	32,596	128
m889	Multi-timescale molecular simulation study of nanoconfined fluids, nanomechanics and self-assembly at organo-metallic interfaces	Yongsheng Leng, George Washington University	Materials Sciences	35,995	6	46,865	20

6. Case Studies

6.1 Chemistry

My reason for using Jacquard is that it is the only platform at NERSC on which I can carry out my calculations.

I do electronic structure calculations on rather small molecules, usually involving 5-8 atoms; for example CH4, C2H4, C2H4O, 2(H2O), 3(H2O), that kind of size. For each molecule I need to calculate the electronic energy for many different geometries, several 104 geometries. Each individual calculation takes, depending on the system and on the precise electronic structure method, anywhere from a minute to an hour; 15 minutes is typical. I do many such calculations in parallel, and it is embarrassingly parallel, with one job per CPU, so two jobs per node on Jacquard.

For the calculations I use the Molpro package almost exclusively, although I might also use Gaussian, Aces, NWChem, or any of a number of other electronic structure packages. To the best of my knowledge, whatever package I use, the hardware issues are the same and I could only use Jacquard at NERSC to make progress. The critical hardware issue is that these calculations require disk space for scratch use. I believe it has to be that way, it is a common feature of electronic structure codes; in any, case, the use of disk space is outside my control.

The electronic structure calculation involves two main phases. The first phase is called integral evaluation, and this is where the scratch disk space is involved. The second phase is an iterative solution of large-scale nonlinear equations. I believe that the second phase does not involve disk.

On Jacquard, by special permission of the consultants, I have set things up so that Molpro uses local disk on each node, so the /tmp filesystem. On Franklin this is not available. Without use of local disk my batch of Molpro jobs, even if it is just 32 of them, completely overwhelms the central filesystem, so /scratch on Jacquard, and my calculations grind to a halt. Worse, of course, my jobs grind the /scratch filesystem to a halt for all users. With use of local disk I can run as many jobs in parallel as the scheduler will give to me. (Even so, the small size of /tmp on Jacquard has been a problem for me that prevents me from doing some of my larger calculations at NERSC.)

I do not know any way around this reliance on /tmp, and I hope that the cluster that is now being procured will offer local disk. Certainly the inner workings of Molpro, Gaussian, NWChem, etc., are beyond my control; and in any case, the use of scratch disk space in those codes is for good reasons, as can be seen in that it is a common feature of all such codes.

I often use Jacquard to run direct ab initio molecular dynamics (AIMD) for studying the dynamics and kinetics of radical reactions. In the calculations, we use the home-developed code (DualOrthGT) that calls a quantum chemistry program such as Gaussian (or MolPro). Both DualOrthGT and Gaussian are well parallelized via MPI. Although Franklin provides the same features as Jacquard, Franklin like IBM Blue Gen does not support "Call SYSTEM()" which is the key function to build a link between MD and ab initio codes. Thus, it's rare to use Franklin. Thank you very much for your kindly supports.

There are several reasons for us to run on Jacquard:

1. Jacquard has more scientific software ready to be used, in particular Gaussian 03. Bassi also has it, but the waiting time for the job to run is much longer.
2. Related to the above, Jacquard has long (48-hour) queues, which is ideal for running programs like G03 (fewer nodes but longer running time). I don't think there is a substitute (yet) for Jacquard for this purpose. These Gaussian jobs are too demanding for us to run on local computers (8-16 processors parallel for each job), but are very important for our research on large organic/inorganic molecules. Without Jacquard, a lot of such research could not have been done.
3. Jacquard has more usable memory per node that Franklin.
4. Most times Jacquard is more stable than Franklin.

I use Jacquard exclusively to run Gaussian 03 jobs of small-medium size in parallel mode using Linda. So, the availability and stability of G03 and Linda on Jacquard is the most important factor. The system seems to be more stable than Franklin and more convinient than Bassi. Queue seems to move faster than on Franklin and admits more tasks than Bassi.

I use Gaussian 03, a computational chemistry software package, for the majority of my calculations at NERSC. There are two reasons I like to use Jacquard instead of Bassi (Gaussian is not available on Franklin, Ed.)

1. Jacquard has a shorter wait time.
2. Gaussian 03 does not parallelize well on Bassi. My personal benchmark has shown that jobs parallelize well, under reasonable scaling factor, up to 16 nodes on Jacquard, but scale very badly if one tries to run multi-node G03 calculations on Bassi.

I saw the announcement of replacing Jacquard by Cray's machine. I just worry that I will have to run all of my Gaussian jobs on Bassi (which is planned for retirement January 2010, Ed.), which is limited to one node, once Jacquard is retired. Since Gaussian is one of the most popular quantum chemistry packages in the computational chemistry community, I really hope NERSC can work with Gaussian Inc. together to figure out a way of putting well-parallel Gaussian program on the new cluster.

For me, the most efficient way to use NERSC resources is the following:

Franklin: Urgent jobs that would take a long time on other machines; the reason is that Franklin is fast, and the line to run is short.

Bassi: long jobs that I need soon but are not urgent; the line to run is longer, but you are allowed to run longer simulations than on Franklin.

Jacquard: some simple calculations that don't require too much simulation time and can be done while urgent jobs are running on Franklin or Bassi. The line to run is short.

For example: I have to run VASP to calculate the adsorption energy of oxygen on a Pt surface, and I need to do DFT and ab-initio molecular dynamics; in addition, once the surface is relaxed I need the atoms' charges. Thus, in order to calculate the adsorption energy is necessary to run one Pt surface, one oxygen atom, and one system with both the oxygen and the surface, all these doing DFT and ab-initio molecular dynamics; moreover, for each system. I have to make an additional short simulation to calculate just the atom charges. So my plan will be: DFT relaxation for each system on Franklin (they are fast on this machine); ab-initio molecular dynamics relaxation for each system in Bassi (they take time, they will run probably on Sunday), and as soon as one of the systems is relaxed I run on Jacquard a simple simulation to calculate the atoms charges.

The idea of my plan is that at some point I'm running all what I need at the same time; therefore, I will get results soon.

6.2 Materials Sciences

I have moved some applications to Franklin. I kept most of my jobs on Jacquard because the charge factor is lower and I usually use fewer than 64 processors. I run large jobs on Franklin, it works great!

We mainly use GROMACS to run Molecular dynamic simulations. We found that GROMACS does not parallelize well beyond 30-40 cores. So, I used Jacquard to run longer simulations (up to 48 hrs) using 32 cores whereas on Franklin I can only run jobs for 24 hrs. That is one of the main advantages.

Secondly, GROMACS was available on Jacquard (I assume that it was compiled to get best efficiency) whereas on Franklin I compiled GROMACS myself and I am not sure how efficient it is.

I can run on Franklin, but the charge rate is lower on Jacquard. The jobs I run are small enough that jacquard is suitable. I run VASP with usually >24 processors. My larger runs are done on Franklin.

The most important things for me are: throughput (mostly, how long is the queue time), availability of VASP, and charge factor.

We use an in-house FORTRAN parallel code to run many jobs investigating the hydration mechanism of different metal ions. We prefer Jacquard because it is about 25% faster than Franklin. This was observed in early 2007 and I had substantial discussion with NERSC staff members. Unfortunately, I did not get very clear explanations. I have not run jobs in Franklin since early 2008 and don't know its current performance. Also, Jacquard is cheaper than Franklin regarding computing time charge.

6.3 Fusion Energy Research

I use Jacquard for small jobs (usually 32 processors) because these jobs have a more rapid turnaround time on Jacquard than on Franklin. I use Franklin for jobs requiring in excess of 100 processors.

Much of the work done on GYRO is done with 32 or 64 processors, but often requires long run times (more than 24 or 48 hours). We generally accomplish this by saving an intermediate phase and resubmitting the job. Naturally, this strategy is facilitated by short wait times in the queue. Emphasis in the PBS algorithm on running large jobs restricts our ability to use the full NERSC capacity. For the 32 processor jobs I run on Jacquard, my needs are well served by increased capacity at the expense of capability. The replacement for Jacquard should emphasize rapid turnaround on small jobs, with a low job size cap (around 100 processors or less). For these needs, a large number of small machines governed by a single batch system would be more than satisfactory. Let users choose between the high capability option (Franklin) and high capacity with low capability.

The other reason I used Jacquard, particularly earlier in the year, was that it was much more stable than Franklin.

The main reason I have been using Jacquard is that I run a lot of small 32 processor jobs. Jacquard has been more reliable and wait time shorter than for Franklin.

6.4 Geosciences

I think there are four main reasons that we use Jacquard as opposed to Bassi, Franklin or DaVinci.

1. We have already done code development, debugging, verification and large scale runs on it;
2. It has more memory per core than Franklin (or Hopper!!!)
3. The two-cores per processor makes it easy for our run to use the entire node. From other clusters, we have experience that the code performance can vary greatly based on which other processes are running on the same node.
4. Machine charge factor is only 60% of Bassi's.